

CLAIMS

What is claimed is:

1 1. A method of modifying data in a database system, the method comprising the steps
2 of:
3 constructing work granules that manipulate rows in a manner that groups the rows
4 within said work granules according to logical storage units that contain the
5 rows; and
6 during execution by an entity of a particular work granule that involves manipulation
7 operations for rows in a particular logical storage unit:
8 causing said entity to perform said manipulation operations for rows
9 completely contained in said logical storage unit;
10 determining a set of spanning rows that are partially contained in said logical
11 storage unit and that satisfy a particular condition; and
12 causing said entity to perform said manipulation operations for all pieces of all
13 spanning rows in said set of spanning rows.

1 2. The method of Claim 1, wherein said particular condition is that each spanning row in
2 said set start in said logical storage unit.

1 3. The method of Claim 1, wherein:
2 said work granules include:
3 a first work granule that involves manipulation operations for a first logical
4 storage unit that includes a portion of a row, and
5 a second work granule that involves manipulation operations for a second
6 logical storage unit that also includes another portion of said row;
7 the method further includes:

8 during execution of said first work granule,
9 determining that said row satisfies said particular condition, and
10 in response to determining that said row satisfies said particular
11 condition, performing a manipulation operation for said row;
12 and
13 during execution of said second work granule,
14 determining that said row does not satisfy said particular condition,
15 and
16 in response to determining that said row does not satisfy said particular
17 condition, foregoing any manipulation operation for said row.

1 4. The method of Claim 1, wherein said work granules are executed as part of a
2 transaction that includes a first subtransaction and a second subtransaction, wherein a
3 first work granule is executed as part of said first subtransaction and a second work
4 granule is executed as part of said second subtransaction.

1 5. The method of Claim 4, wherein
2 said first work granule involves manipulation operations for a first logical storage unit
3 that includes a portion of a row;
4 said second work granule involves manipulation operations for a second logical
5 storage unit that also includes another portion of said row; and
6 the method further includes executing a manipulation operation for said row as part of
7 said first subtransaction.

1 6. The method of Claim 5, wherein said first logical storage unit contains a first
2 transaction list with a first entry, wherein said second logical storage unit contains a
3 second transaction list with a second entry, wherein said step of executing a

manipulation operation for said row includes assigning ownership of said first entry and said second entry to said first subtransaction.

7. The method of Claim 1, wherein said particular logical storage unit includes a data block in said database system.

8. A method of inserting rows into logical storage units that store information in a database system, the method comprising the steps of:
inserting a first row piece of a spanning row into a first logical storage unit;
prior to inserting a second row piece of said spanning row into a second logical storage unit, determining whether one or more criteria is satisfied, wherein said one or more criteria include that said second logical storage unit has enough space allocated to identify at least a threshold number of interested transactions; and
inserting said second row piece of said spanning row into said second logical storage unit only when said one or more criteria are satisfied.

9. The method of Claim 8, wherein a data structure in each logical storage unit of said logical storage units is used to identify interested transactions for said each logical storage unit.

10. The method of Claim 9, wherein said data structure is a transaction list with entries, wherein each entry of said entries may be owned by an interested transaction.

11. The method of Claim 8, wherein said one or more criteria include that said second row piece be the second or greater row piece in said spanning row.

1 12. The method of Claim 8, wherein said threshold number is greater than the sum of the
2 quantity of overflow row pieces stored in said second logical storage unit after
3 inserting said second row piece.

1 13. A computer-readable medium carrying one or more sequences of instructions for
2 modifying data in a database system, wherein execution of the one or more sequences
3 of instructions by one or more processors causes the one or more processors to
4 perform the steps of:
5 constructing work granules that manipulate rows in a manner that groups the rows
6 within said work granules according logical storage units that contain the
7 rows; and
8 during execution by an entity of a particular work granule that involves manipulation
9 operations for rows in a particular logical storage unit:
10 causing said entity to perform said manipulation operations for rows
11 completely contained in said logical storage unit;
12 determining a set of spanning rows that are partially contained in said logical
13 storage unit and that satisfy a particular condition; and
14 causing said entity to perform said manipulation operations for all pieces of all
15 spanning rows in said set of spanning rows.

1 14. The computer-readable media of Claim 13, wherein said particular condition is that
2 each spanning row in said set start in said logical storage unit.

1 15. The computer-readable media of Claim 13, wherein:
2 said work granules include:
3 a first work granule that involves manipulation operations for a first logical
4 storage unit that includes a portion of a row, and

5 a second work granule that involves manipulation operations for a second
6 logical storage unit that also includes another portion of said row;
7 the computer-readable media further includes instructions for performing:
8 during execution of said first work granule,
9 determining that said row satisfies said particular condition, and
10 in response to determining that said row satisfies said particular
11 condition, performing a manipulation operation for said row;
12 and
13 during execution of said second work granule,
14 determining that said row does not satisfy said particular condition,
15 and
16 in response to determining that said row does not satisfy said particular
17 condition, foregoing any manipulation operation for said row.

1 16. The computer-readable media of Claim 13, wherein said work granules are executed
2 as part of a transaction that includes a first subtransaction and a second
3 subtransaction, wherein a first work granule is executed as part of said first
4 subtransaction and a second work granule is executed as part of said second
5 subtransaction.

1 17. The computer-readable media of Claim 16, wherein
2 said first work granule involves manipulation operations for a first logical storage unit
3 that includes a portion of a row;
4 said second work granule involves manipulation operations for a second logical
5 storage unit that also includes another portion of said row; and

the computer-readable media further includes instructions for executing a manipulation operation for said row as part of said first subtransaction.

18. The computer-readable media of Claim 17, wherein said first logical storage unit contains a first transaction list with a first entry, wherein said second logical storage unit contains a second transaction list with a second entry, wherein said step of executing a manipulation operation for said row includes assigning ownership of said first entry and said second entry to said first subtransaction.

19. The computer-readable media of Claim 13, wherein said particular logical storage unit includes a data block in said database system.

20. A computer-readable medium carrying one or more sequences of instructions for inserting rows into logical storage units that store information in a database system, wherein execution of the one or more sequences of instructions by one or more processors causes the one or more processors to perform the steps of:
inserting a first row piece of a spanning row into a first logical storage unit;
prior to inserting a second row piece of said spanning row into a second logical storage unit, determining whether one or more criteria is satisfied, wherein said one or more criteria include that said second logical storage unit has enough space allocated to identify at least a threshold number of interested transactions; and
inserting said second row piece of said spanning row into said second logical storage unit only when said one or more criteria are satisfied.

1 21. The computer-readable media of Claim 20, wherein a data structure in each logical
2 storage unit of said logical storage units is used to identify interested transactions for
3 said each logical storage unit.

1 22. The computer-readable media of Claim 21, wherein said data structure is a
2 transaction list with entries, wherein each entry of said entries may be owned by an
3 interested transaction.

1 23. The computer-readable media of Claim 20, wherein said one or more criteria include
2 that said second row piece be the second or greater row piece in said spanning row.

1 24. The computer-readable media of Claim 20, wherein said threshold number is greater
2 than the sum of the quantity of overflow row pieces stored in said second logical
3 storage unit after inserting said second row piece.